

12/23/99
JC662 U.S. PT

Frederick P. Fish
1855-1930

W.K. Richardson
1859-1951

12-2799
FISH & RICHARDSON P.C.
December 23, 1999

4225 Executive Square
Suite 1400
La Jolla, California
92037

Telephone
858 678-5070

Faximile
858 678-5099

Web Site
www.fr.com

12/23/99
JC662 U.S. PT
09/41364

Attorney Docket No.: 10559/096001/P7615

Box Patent Application

Assistant Commissioner for Patents
Washington, DC 20231

Presented for filing is a new original patent application of:

Applicant: CYNTHIA L. BICKERSTAFF, MARK A. VANANTWERP, WILLIAM W. LOVE, STACY P. PURCELL, LAWRENCE J. WIDMER AND JEFFREY C. SEDAYAO

Title: REAL-TIME PERFORMANCE ASSESSMENT OF LARGE AREA NETWORK USER EXPERIENCE

BOSTON
DELAWARE
NEW YORK
SILICON VALLEY
SOUTHERN CALIFORNIA
TWIN CITIES
WASHINGTON, DC

Enclosed are the following papers, including those required to receive a filing date under 37 CFR 1.53(b):

Pages

Specification	19
Claims	12
Abstract	1
Declaration	[To be Filed at a Later Date]
Drawing(s)	4

Enclosures:

— Postcard.

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL528179594US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit December 23, 1999

Signature

R. Donovan

Typed or Printed Name of Person Signing Certificate

FISH & RICHARDSON P.C.

Assistant Commissioner for Patents

December 23, 1999

Page 2

Basic filing fee	\$0
Total claims in excess of 20 times \$18	\$0
Independent claims in excess of 3 times \$78	\$0
Fee for multiple dependent claims	\$0
Total filing fee:	\$0

No filing fee is being paid at this time. Please apply any other required fees, **EXCEPT FOR THE FILING FEE**, to deposit account 06-1050, referencing the attorney document number shown above. A duplicate copy of this transmittal letter is attached.

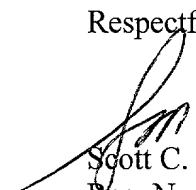
If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (858) 678-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

SCOTT C. HARRIS
Fish & Richardson P.C.
4225 Executive Square, Suite 1400
La Jolla, CA 92037

Respectfully submitted,



Scott C. Harris
Reg. No. 32,030
Enclosures
SCH/jzc
10012260.doc

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: REAL-TIME PERFORMANCE ASSESSMENT OF LARGE AREA NETWORK USER EXPERIENCE

APPLICANT: CYNTHIA L. BICKERSTAFF, MARK A. VANANTWERP, WILLIAM W. LOVE, STACY P. PURCELL, LAWRENCE J. WIDMER AND JEFFREY C. SEDAYAO

CERTIFICATE OF MAILING BY EXPRESS MAIL

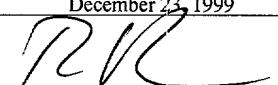
Express Mail Label No. EL528179594US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

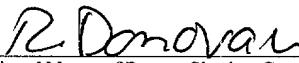
Date of Deposit

December 23, 1999

Signature



Typed or Printed Name of Person Signing Certificate



**REAL-TIME PERFORMANCE ASSESSMENT OF
LARGE AREA NETWORK USER EXPERIENCE**

TECHNICAL FIELD

This invention relates to measuring and modifying the performance of a network, and more particularly to a method, system, and computer program for real-time measurement and modification of the performance of communications on a large area network, such as the Internet, based upon actual user experience.

BACKGROUND

FIG. 1 shows a typical configuration of a large area network, such as the Internet. Multiple geographically dispersed user client systems 100 are connected through user Internet Service Providers (ISPs) 102 to the network "cloud" 104 comprising the Internet backbone communication transport systems. Requests for information or services by a client application program executing on a user client system 100 may be routed through multiple server ISPs 106 through a router 108 to a web server 110. The web server 110 retrieves or generates the requested information or provides the requested service, and communicates a response to the requesting client application.

Various attempts have been made to actively assess the performance (e.g., response time, transmission problems, etc.)

of Internet connections. Some current methods use active measurement of various network or server components that are dedicated to performance measurement. For example, conventional Internet Control Message Protocol (ICMP) "ping" and "traceroute" commands can be used to measure the performance of the network connections between a client terminal and a server. However, these commands are frequently transmitted with a different (often lower) priority than the protocols used by applications run by users for "web surfing".

10 Accordingly, inaccurate (*i.e.*, false positive) measurements are common.

Attempts have been made to actively measure Internet connections using the same protocols used by end-user applications, such as the HyperText Transfer Protocol (HTTP) "GET" command. These approaches typically use computer programs (sometimes known as "hosts", "agents", or "beacons") residing on measurement instrumentation capable of communicating with Internet protocols. However, such computer programs are limited to assessing network paths only from the specific network nodes on which they are executing. Further these techniques inject traffic into sometimes overburdened Internet, WAN, or LAN infrastructures, causing the measurement process to change the characteristics being measured. Additionally, these techniques are relatively expensive to implement.

A further problem of all of these active or injected measurement approaches is that they generate non-value added communication traffic for both local and large area network infrastructures.

5

SUMMARY

In one aspect, the invention includes a method, system, and computer program for real-time measurement of the performance of communications on a large area network between a selected server and a plurality of users, based upon actual user experience, including: accessing a server log having records of actual user access to the selected server; aggregating records from the server log into a plurality of aggregate slots, each having at least one time bin, based on an aggregation method; performing at least one statistical analysis of each time bin of each aggregate slot; and outputting the results of such statistical analysis as an indication of actual server usage by users.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 shows a typical configuration of a large area network, such as the Internet.

FIG. 2 is a process flow diagram showing one embodiment of the invention.

5 FIG. 3 is a flowchart for a process comparing information from a Classless Inter-Domain Routing (CIDR) block database and an Internet Protocol (IP) address input in order to convert the IP address to geographic or source information according to one embodiment of the invention.

10 FIG. 4 is a flowchart showing an embodiment for modifying traffic paths through a router to the Internet.

Like reference numbers and designations in the various drawings indicate like elements.

15 **DETAILED DESCRIPTION**

Embodiments of the invention are directed to a method, system, and computer program for real-time measurement and modification of the performance of communications on large area networks, such as the Internet, based upon actual user experience. One embodiment performs a statistical analysis of access logs that record actual server usage by users. Based on such analysis, routing of communications over the network can be modified to improve overall communications performance.

20 Embodiments may also output results indicative of overall communications performance and of server applications that

interact poorly or especially well with network conditions, thus providing direction to application development efforts.

More particularly, one embodiment of the invention creates correlation assessments of performance related 5 measurements against the geographical location of and/or routes taken by client applications. A route is determined by aggregating client Internet Protocol (IP) addresses according to Classless Inter-Domain Routing (CIDR) blocks or route advertisements available in a conventional fashion by querying 10 a router or router server for such advertisements.

The results of these analyses define which geographical location or route may be performing better or worse than a comparative geographical location or route. Based on such comparisons, active steps may be taken to modify routing of 15 network traffic to increase overall client-server performance.

In addition, each web server running a set of applications can be compared with every other server running a set of applications within the same domain. Such a comparison can detect differences in configuration of the servers, and 20 permits identification of servers that are providing poor performance to users. Based on such comparisons, active steps may be taken to modify the configurations of dissimilar servers to match the performance of other servers within a group of evaluated servers.

An advantage of using such log file analysis over active measurements for detecting performance over a large area network such as the Internet is that historical records of an end user's experience can be mined for objective quantitative information and compared to the experience of other end users collected at or near the same time. This allows for identifying the root cause of performance problems. Since actual user experience is assessed, the limitation of a few, expensive sampling locations of beacons or agents is eliminated. The integrity of the analysis for any individual web site is limited by the popularity of a web site residing on a web server. However, an enterprise hosting multiple web sites can alleviate this limitation by aggregating across multiple web sites for the same end-user population. Tuning the performance of a web site to those users already using it enhances current users' experience.

One embodiment of the invention creates Pareto analyses of different applications running on a server where applications taking longer than a configurable time interval for greater percentages of the use of the application by end users are sorted in order from "most often" to "least often" exceeding the interval. Based on such an analysis, allocation of application developer resources to poorest performing applications can be made to improve the application and improve the end user experience.

Statistical Analysis Process Flow

FIG. 2 is a process flow diagram showing an embodiment of the invention adapted for use with the Internet. One input to the process is a server log file 200 that is maintained on a web server 110 (see FIG. 1) or a mounted file system in conventional fashion. In the illustrated embodiment, the server log file 200 is configured according to World Wide Web Consortium (W3C) standards, and includes a log file 202 for each web server 110 being monitored within a server group. A typical server log file 200 configured to such standards records all user accesses of every element of a web page. A typical logical organization for the server log file 200 is a table having columns for every recorded data item, and rows for each access event. In particular, a server log file 200 that is most suitable for use with the invention records the following data for each user access of every element of a monitored web page: a time stamp indicating when a record was created in the server log file 200; client IP (c_ip) address; bytes transmitted from client application to server (cs_bytes); bytes transmitted from server to client application (sc_bytes); the time taken to complete a two-way transmission of bytes between a client and a server

(time_taken); status codes documenting an action for each web page element (status_code); each URL (uniform resource locator) requested by a client browser (uniform resource identifier stem, or uri_stem); the type of browser used by the 5 end-user (user-agent); and each URL referring to the uri_stem (uniform resource identifier referrer, or uri_referrer). Other information may also be included in each record, as desired or in order to comply with Internet standards.

In the illustrated embodiment, any particular server log 10 file 200 is closed to new data entries before commencement of any statistical analysis. A new log file may be opened in known fashion to continue to record user access while the closed log file 200 is analyzed. Embodiments of the invention may also use log file entries written from a server directly 15 to a flat file or database.

Process parameters are defined in a process settings step 202. In this step, an analyst either selects an aggregation method (e.g., "aggregate by log-file column", or "aggregate by client IP address"), optional filtration parameters, and an 20 aggregation bin time increment, or these parameters set by reference to default (i.e., pre-established) settings.

If filtration parameters are set in the process settings step 202, the data in the server log file 200 is filtered to remove records that are not to be counted in further 25 statistical analyses. For example, such records may be from

non-customer sources, such as a beacon or agent, and thus do not reflect actual user accesses to the web server 110. In the illustrated embodiment, an agent ID field within a conventional W3C compliant server log file is used to filter out undesirable records. However, any desired record field may be used to perform a selected filtration. In the illustrated embodiment, filtering is implemented as a string matching function that compares a filter string to any character string or substring in any of the log file fields. Other types of filtering may be employed, such as by comparing the client IP (c_ip) address against a "lookup" table of addresses to include or exclude.

The selected log file records are then processed in an aggregation step 208 using the aggregation method defined in the process settings step 202. Typical aggregation methods are an "aggregate by log-file column" method 210 (e.g., AS-path, country, region, etc.) or an "aggregate by client IP address" method 212. The aggregation method creates entries within an aggregation table 216 having multiple aggregate slots 218 each generally having multiple time bins 220.

For example, the log-file column aggregation method 210 reads a defined record column (or "field") data value and time stamp for each selected log file record and assigns that record to a corresponding time bin 220 within the appropriate aggregate slot 218. Thus, records accumulated over a 24-hour

period and corresponding to a first defined column data value can be assigned to 24 1-hour time bins 220 in a first aggregate slot 218, while records corresponding to a second defined column data value are assigned to 24 1-hour time bins 5 220 in a second aggregate slot 218.

If the "aggregate by client IP address" method 212 is selected in the process settings step 202, it is generally desirable to convert the raw IP address of a user client system 100 accessing the web server 110 to a geographic 10 location or specific source (e.g., country, region, company, and/or ISP). In the illustrated embodiment, this is accomplished by supplying the client IP (c_ip) address from each record to an IP Lookup function 214, which returns geographic location or specific source information associated 15 with that address. One implementation of the IP Lookup function 214 is described in detail below.

Once the aggregation table 216 has been created, the log file entries within each time bin 220' may be subjected to a statistical analysis process 222. This process applies a 20 variety of statistical analysis algorithms 224 to derive information on server usage and statistical significance of such information based on the actual user access records. Collections of multiples of such time-bins 220' can also be assembled in chronological order to determine trends for each 25 of the statistical measures. In the illustrated embodiment,

the basic rate and count information computed are: byte-density (computed as sc_bytes + cs_bytes); transfer-rate (computed as byte-density divided by time taken); URL-count (total number of log entries); error-fraction (the fraction of all log entries having errors); cache-fraction (the fraction of all log entries having cached URL's as determined by response code); and unique-IP-address-count (the number of unique IP addresses among all log entries). Once the basic rate and count information is computed, distribution statistics may be computed for some or all of such basic information. In particular, in the illustrated embodiment, distribution statistics, such as quartiles, interquartile range (IQR), and median, are computed in known fashion for the byte-density and transfer-rate statistics.

The results of the statistical analysis process 222 can generate output 226 in several forms. The raw data from the statistical analysis process 222 can be output directly. Trend information 228 can be output (e.g., in table or graphical form) to show the trends of time bins by aggregate slot or item. Various comparison tests (e.g., out of range, over threshold, percentage change, etc.) can be applied to the basic rate and count information as well as the distribution statistics to trigger an event notification 230 (e.g., notice to a network administrator) if any selected statistical value is abnormal. Further, the statistical information for multiple

time bins and/or aggregate items can be input to various comparison tools 232 for troubleshooting. For example, the IQR statistics for the byte-density for two servers within a domain can be compared graphically for visual assessment by a network administrator. Generation of such trend displays, event notifications, and comparisons is well-known in the art.

Thus, the illustrated embodiment of the invention can create correlation assessments of performance related measurements against the geographical location and/or route traversed during use of a network application by an end-user.

In particular, transfer-rate and error-fraction measurements can be correlated to at least the following parameters: geographical location of c_ip addresses; ISP for c_ip addresses; net block or route of c_ip addresses; and applications requested (uri_stem) or previously requested (uri_referrer) by client applications or users from the web server 110. The results of these analyses define which geographical location, ISP, net block, route, or application may be performing better or worse than a comparative geographical location, ISP, net block, route, or application.

The validity of the correlations is ensured by performing statistical validity checks between applications and servers, such as by ensuring similarity or sufficiency of certain of the computer distribution statistics, in known fashion. The byte-density, URL-count, and unique-IP-address-count

parameters are used to ensure valid correlations. For example, since the common TCP/IP protocol (the protocol used over the Internet) changes its performance based on the number of packets transmitted (through congestion control and "slow start" mechanisms), requiring a similar value for the byte-density parameter ensures that differences between servers or services of different applications are due to other interesting parameters (such as the geographical location of c_ip addresses, ISP for c_ip addresses; etc.) instead of resulting from artifacts (e.g., large byte transfers generated by the TCP/IP protocol itself). The combination of the URL-count and the unique-IP-address-count parameters represent the sample size of the analysis space. Since each unique IP address essentially represents a different end-to-end communication path, the unique-IP-address-count measures the diversity of the network space being measured. Requiring that the URL-count and the unique-IP-address-count parameters exceed a selected threshold helps ensures that the correlations described above are valid.

If the correlations described above indicate a problem, actions may be undertaken to rectify the problem. These actions may include: selecting a better exit path from a multi-homed (*i.e.*, having multiple ISPs) data center (described in greater detail below); notifying a network administrator to repair a server which is performing below the

level of ostensibly identical servers; and indicating the need to re-write applications which are slow performing. For this later case, it may be that the applications perform well during local area network testing, but log file analysis in accordance with the invention may reveal an application specific sensitivity to actual Internet conditions.

IP Lookup function

FIG. 3 is a flowchart showing an embodiment for comparing information from a Classless Inter-Domain Routing (CIDR) block database 300 and an IP address input 304 in order to convert the IP address to geographic or source information. A CIDR block defines a subnet of a larger network. A CIDR address includes a standard 32-bit IP address and also information on how many bits are used for the network prefix. This addressing scheme allows for efficient allocation of IP addresses compared to prior standards. The CIDR addressing scheme also enables "route aggregation" in which a single high-level route entry can represent many lower-level routes in global routing tables.

In the illustrated embodiment for the Internet, the CIDR block database 300 is specially generated by querying (using conventional Internet query commands) regional Internet registries for CIDR blocks that have been assigned through such registries. The responses from the registries include

CIDR block address (used as the database key), country code, network name, network description, region (*i.e.*, sub-country geographical information, sometimes down to a street address), and date of the last update for each registry record.

5 During operations, the CIDR block database 300 may be read into memory organized as a 32-element array 302. Each array element 303 is a binary tree of CIDR block records selected with a unique subnet mask value. For example, array element "0" contains a binary tree of all CIDR block records whose subnet mask is "255.255.255.255" (*i.e.*, having a binary representation of 32 "1's"). Similarly, array element "1" contains a binary tree of all CIDR block records whose subnet mask is "255.255.255.254" (*i.e.*, having a binary representation of 31 "1's" followed by one "0" as the least significant bit). This pattern continues, such that array element "31" contains a binary tree of all CIDR block records whose subnet mask is "1.0.0.0" (*i.e.*, having a binary representation of one "1" followed by 31 "0's").

10 The subnet mask for each array element is used to mask a target IP address before searching the element's associated binary tree. The subnet mask can be computed from the CIDR block mask number as the binary complement of $2^{32-\text{MaskNumber}} - 1$. This configuration of CIDR blocks in the memory array 302 provides for most specific CIDR block/IP address matching.

In operation, a target c_ip address from a record in the server log 200 is used as input to the most specific CIDR block/IP address matching process (STEP 304). For each c_ip address, a counter N is set to "0", representing array element 5 "0" (STEP 306). Using the subnet masking technique described above, the target c_ip address is masked with the array element's associated subnet mask (e.g., all "1's" for array element "0"), and the corresponding array element's binary tree is then traversed to find a record match (STEP 308). In 10 particular, the masked IP address component of each CIDR block for each record traversed in the Nth binary tree is compared against the masked target IP address.

If a match exists, then desired record fields from the corresponding CIDR block (e.g., country code, network name, 15 network description, region, and/or date) are sent to output to be used for binning by the lookup requestor (STEP 310). Thus, the c_ip address is converted to geographical and source information.

If no match occurs (STEP 308), N is incremented and 20 tested for being in the range 0-31 (STEP 312). If N is out of range, no match exists and is so indicated (STEP 314). Otherwise, the match process continues with the next array element through similar masking of the target c_ip address and traversal of the associated binary tree for the incremented 25 value of N.

Active ISP Routing

FIG. 4 is a flowchart showing an embodiment for modifying traffic paths through a router to a large area network such as the Internet. After the statistical analysis described above is performed, the results can be used to "tune" performance of a server system. In the illustrated embodiment, the exit route for communications from a web server 110 through the router 108 and all connected server ISPs 106 to the network 104 is determined for each c_ip address (STEP 400). This may be accomplished by querying the router 108 (or a route server), using conventional network control commands, for the routing table maintained by the router 108. The routing information may then be analyzed to determine which exit path has the highest performance (e.g., highest transfer-rate for a particular destination).

Once a preferred exit route is determined, the routing of traffic may be biased towards that exit route (or, alternatively, away from the most poorly performing exit routes). For the Internet, this may be done using Border Gateway Protocol (BGP) mechanisms. BGP is commonly used as a router-to-router protocol between administrative domains. For example, in the illustrated embodiment, outgoing traffic is biased by modifying incoming routing update information using BGP path prepending or local preference mechanisms. Similarly, incoming traffic is biased by modifying outgoing routing

update information using BGP path prepending or community string mechanisms.

Implementation

The invention may be implemented in hardware or software,
5 or a combination of both (e.g., programmable logic arrays).

Unless otherwise specified, the algorithms included as part of
the invention are not inherently related to any particular
computer or other apparatus. In particular, various general
purpose machines may be used with programs written in

10 accordance with the teachings herein, or it may be more
convenient to construct more specialized apparatus to perform
the required method steps. However, preferably, the invention
is implemented in one or more computer programs executing on
programmable systems each comprising at least one processor,
15 at least one data storage system (including volatile and non-
volatile memory and/or storage elements), at least one input
device, and at least one output device. The program code is
executed on the processors to perform the functions described
above.

20 Each such program may be implemented in any desired
computer language (including machine, assembly, or high level
procedural, logical, or object oriented programming languages)
to communicate with a computer system. In any case, the
language may be a compiled or interpreted language.

Each such computer program may be stored on a storage media or device (e.g., solid state, magnetic, or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the
5 storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a
10 specific and predefined manner to perform the functions described herein.

A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other
15 embodiments are within the scope of the following claims.

WHAT IS CLAIMED IS:

1 1. A method for real-time measurement of the
2 performance of communications on a large area network between
3 a selected server and a plurality of users, based upon actual
4 user experience, including:

5 (a) accessing a server log having records of actual user
6 access to the selected server;
7 (b) aggregating records from the server log into a plurality
8 of aggregate slots, each having at least one time bin,
9 based on an aggregation method;
10 (c) performing at least one statistical analysis of each
11 time bin of each aggregate slot; and
12 (d) outputting the results of such statistical analysis as
13 an indication of actual server usage by users.

1 2. The method of claim 1, further including filtering
2 out selected records from the server log before the step of
3 aggregating.

1 3. The method of claim 1, further including generating
2 an event notification if a selected statistical analysis value
3 is abnormal.

1 4. The method of claim 1, further including selecting
2 the aggregation method from a set of aggregation methods.

1 5. The method of claim 1, wherein the aggregation
2 method includes aggregation by log-file record column data
3 value for each record from the server log.

1 6. The method of claim 1, further including:
2 (e) determining geographical or source information for each
3 record; and
4 (f) selecting the aggregation method to aggregate records
5 based on such geographical or source information.
6

7 7. The method of claim 6, wherein determining
8 geographical or source information for each record includes:
9 (g) defining a database comprising large area network
10 address blocks having geographical or source
11 information;
12 (h) comparing an address field in each record to the address
13 blocks in the database; and
14 (i) associating with each record the geographical or source
15 information from an address block matching the address
16 field of the record.

17
18 8. The method of claim 7, wherein comparing an address
19 field in each record to the address blocks in the database
20 includes:

21 (j) defining an array of binary trees for the address blocks
22 in the database, each address block within a binary tree
23 within an array element being masked by a corresponding
24 unique subnet mask value;
25 (k) masking each address field in each record by a unique
26 subnet value corresponding to a selected array element;
27 (l) comparing each masked address field to an address field
28 of the address blocks within the binary tree of the
29 selected array element;
30 (m) outputting selected fields of any matching address
31 block; and
32 (n) otherwise, continuing the step of comparing with a next
33 selected array element until a match is found or all
34 array elements have been compared.

1 9. The method of claim 1, further including:
2 (o) determining exit routing paths from each selected server
3 based on the records from the server log;
4 (p) determining a best performing exit route based on the
5 statistical analysis of records from the server log;
6 (q) biasing incoming and outgoing communications with
7 respect to each server to use the determined best
8 performing exit route.
9

10. A method for comparing an address field of a large area network record to a database comprising large area network address blocks having geographical or source information, including:

(r) defining an array of binary trees for the address blocks in the database, each address block within a binary tree within an array element being masked by a corresponding unique subnet mask value;

(s) masking the address field of a large area network record by a unique subnet value corresponding to a selected array element;

(t) comparing each masked address field to an address field of the address blocks within the binary tree of the selected array element;

(u) indicating a match; and

(v) otherwise, continuing the step of comparing with a next selected array element until a match is found or all array elements have been compared.

11. A system for real-time measurement of the performance of communications on a large area network between a selected server and a plurality of users, based upon actual user experience, including:

(w) a server log having records of actual user access to the selected server;

7 (x) means for accessing and aggregating records from the
8 server log into a plurality of aggregate slots, each
9 having at least one time bin, based on an aggregation
10 method;
11 (y) means for performing at least one statistical analysis
12 of each time bin of each aggregate slot; and
13 (z) means for outputting the results of such statistical
14 analysis as an indication of actual server usage by
15 users.

1 12. The system of claim 11, further including means for
2 filtering out selected records from the server log before the
3 step of aggregating.

1 13. The system of claim 11, further including means for
2 generating an event notification if a selected statistical
3 analysis value is abnormal.

1 14. The system of claim 11, further including means for
2 selecting the aggregation method from a set of aggregation
3 methods.

1 15. The system of claim 11, wherein the aggregation
2 method includes aggregation by log-file record column data
3 value for each record from the server log.

1 16. The system of claim 11, further including:

2 (aa) means for determining geographical or source
3 information for each record; and
4 (bb) means for selecting the aggregation method to
5 aggregate records based on such geographical or source
6 information.

1 17. The system of claim 16, wherein the means for
2 determining geographical or source information for each record
3 includes:

4 (cc) a database comprising large area network address
5 blocks having geographical or source information;
6 (dd) a comparison function for comparing an address field
7 in each record to the address blocks in the database;
8 and
9 (ee) an associating function for associating with each
10 record the geographical or source information from an
11 address block matching the address field of the record.

1 18. The system of claim 17, wherein the comparison
2 function includes:

3 (ff) an array of binary trees for the address blocks in
4 the database, each address block within a binary tree
5 within an array element being masked by a corresponding
6 unique subnet mask value;

7 (gg) means for masking each address field in each record
8 by a unique subnet value corresponding to a selected
9 array element;
10 (hh) means for comparing each masked address field to an
11 address field of the address blocks within the binary
12 tree of the selected array element;
13 (ii) means for outputting selected fields of any matching
14 address block; and
15 (jj) means for otherwise continuing the step of comparing
16 with a next selected array element until a match is
17 found or all array elements have been compared.

1 19. The system of claim 11, further including:
2 (kk) means for determining exit routing paths from each
3 selected server based on the records from the server
4 log;
5 (ll) means for determining a best performing exit route
6 based on the statistical analysis of records from the
7 server log;
8 (mm) means for biasing incoming and outgoing
9 communications with respect to each server to use the
10 determined best performing exit route.

1 20. A system for comparing an address field of a large
2 area network record to a database comprising large area

3 network address blocks having geographical or source
4 information, including:
5 (nn) an array of binary trees for the address blocks in
6 the database, each address block within a binary tree
7 within an array element being masked by a corresponding
8 unique subnet mask value;
9 (oo) means for masking the address field of a large area
10 network record by a unique subnet value corresponding to
11 a selected array element;
12 (pp) means for comparing each masked address field to an
13 address field of the address blocks within the binary
14 tree of the selected array element;
15 (qq) means for indicating a match; and
16 (rr) means for otherwise continuing the step of comparing
17 with a next selected array element until a match is
18 found or all array elements have been compared.

1 21. A computer program, stored on a computer-readable
2 medium, for real-time measurement of the performance of
3 communications on a large area network between a selected
4 server and a plurality of users, based upon actual user
5 experience, the computer program comprising instructions
6 for causing a computer system to:

7 (ss) access a server log having records of actual user
8 access to the selected server;

9 (tt) aggregate records from the server log into a
10 plurality of aggregate slots, each having at least one
11 time bin, based on an aggregation method;
12 (uu) perform at least one statistical analysis of each
13 time bin of each aggregate slot; and
14 (vv) output the results of such statistical analysis as
15 an indication of actual server usage by users.

1 22. The computer program of claim 21, further including
2 instructions for causing the computer system to filter out
3 selected records from the server log before the step of
4 aggregating.

1 23. The computer program of claim 21, further including
2 instructions for causing the computer system to generate an
3 event notification if a selected statistical analysis value is
4 abnormal.

1 24. The computer program of claim 21, further including
2 instructions for causing the computer system to select the
3 aggregation method from a set of aggregation methods.

1 25. The computer program of claim 21, wherein the
2 aggregation method includes aggregation by log-file record
3 column data value for each record from the server log.

1 26. The computer program of claim 21, further including
2 instructions for causing the computer system to:

3 (ww) determine geographical or source information for
4 each record; and
5 (xx) select the aggregation method to aggregate records
6 based on such geographical or source information.

1 27. The computer program of claim 26, wherein the
2 instructions for causing the computer system to determine
3 geographical or source information for each record further
4 include instructions for causing the computer system to:

5 (yy) define a database comprising large area network
6 address blocks having geographical or source
7 information;
8 (zz) compare an address field in each record to the
9 address blocks in the database; and
10 (aaa) associate with each record the geographical or
11 source information from an address block matching the
12 address field of the record.

1 28. The computer program of claim 27, wherein the
2 instructions for causing the computer system to compare
3 an address field in each record to the address blocks in
4 the database include instructions for causing the
5 computer system to:

6 (bbb) define an array of binary trees for the address
7 blocks in the database, each address block within a

8 binary tree within an array element being masked by a
9 corresponding unique subnet mask value;
10 (ccc) mask each address field in each record by a unique
11 subnet value corresponding to a selected array element;
12 (ddd) compare each masked address field to an address
13 field of the address blocks within the binary tree of
14 the selected array element;
15 (eee) output selected fields of any matching address
16 block; and
17 fff) otherwise, continue the step of comparing with a
18 next selected array element until a match is found or
19 all array elements have been compared.

1 29. The computer program of claim 21, further including
2 instructions for causing the computer system to:

3 (ggg) determine exit routing paths from each selected
4 server based on the records from the server log;
5 (hhh) determine a best performing exit route based on the
6 statistical analysis of records from the server log;
7 (iii) bias incoming and outgoing communications with
8 respect to each server to use the determined best
9 performing exit route.

1 30. A computer program, stored on a computer-readable
2 medium, for comparing an address field of a large area
3 network record to a database comprising large area

4 network address blocks having geographical or source
5 information, the computer program comprising instructions
6 for causing a computer system to:
7 (jjj) define an array of binary trees for the address
8 blocks in the database, each address block within a
9 binary tree within an array element being masked by a
10 corresponding unique subnet mask value;
11 (kkk) mask the address field of a large area network
12 record by a unique subnet value corresponding to a
13 selected array element;
14 (lll) compare each masked address field to an address
15 field of the address blocks within the binary tree of
16 the selected array element;
17 (mmm) indicate a match; and
18 (nnn) otherwise, continue the step of comparing with a
19 next selected array element until a match is found or
20 all array elements have been compared.

ABSTRACT

A method, system, and computer program for real-time measurement and modification of the performance of communications on a large area network, such as the Internet,
5 based upon actual user experience. One embodiment performs a statistical analysis of access logs that record actual server usage by users. Based on such analysis, routing of communications over the network can be modified to improve overall communications performance.

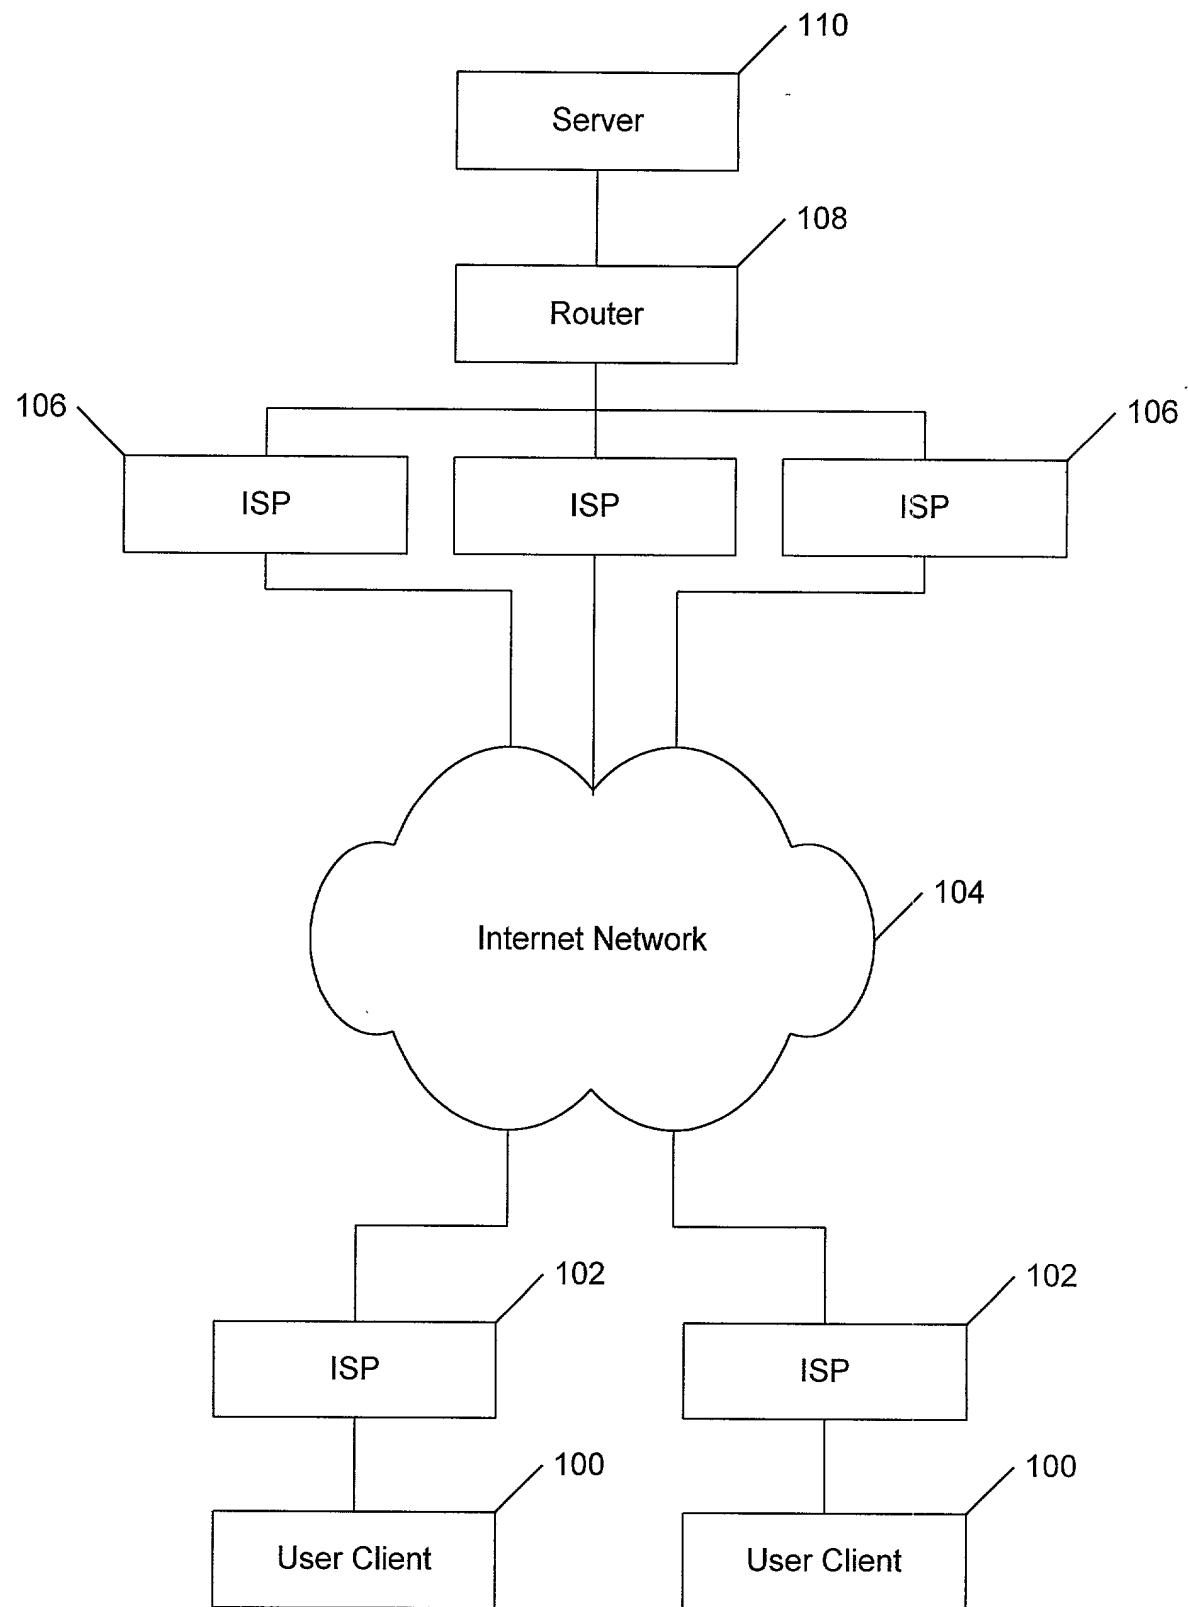


FIG. 1
(prior art)

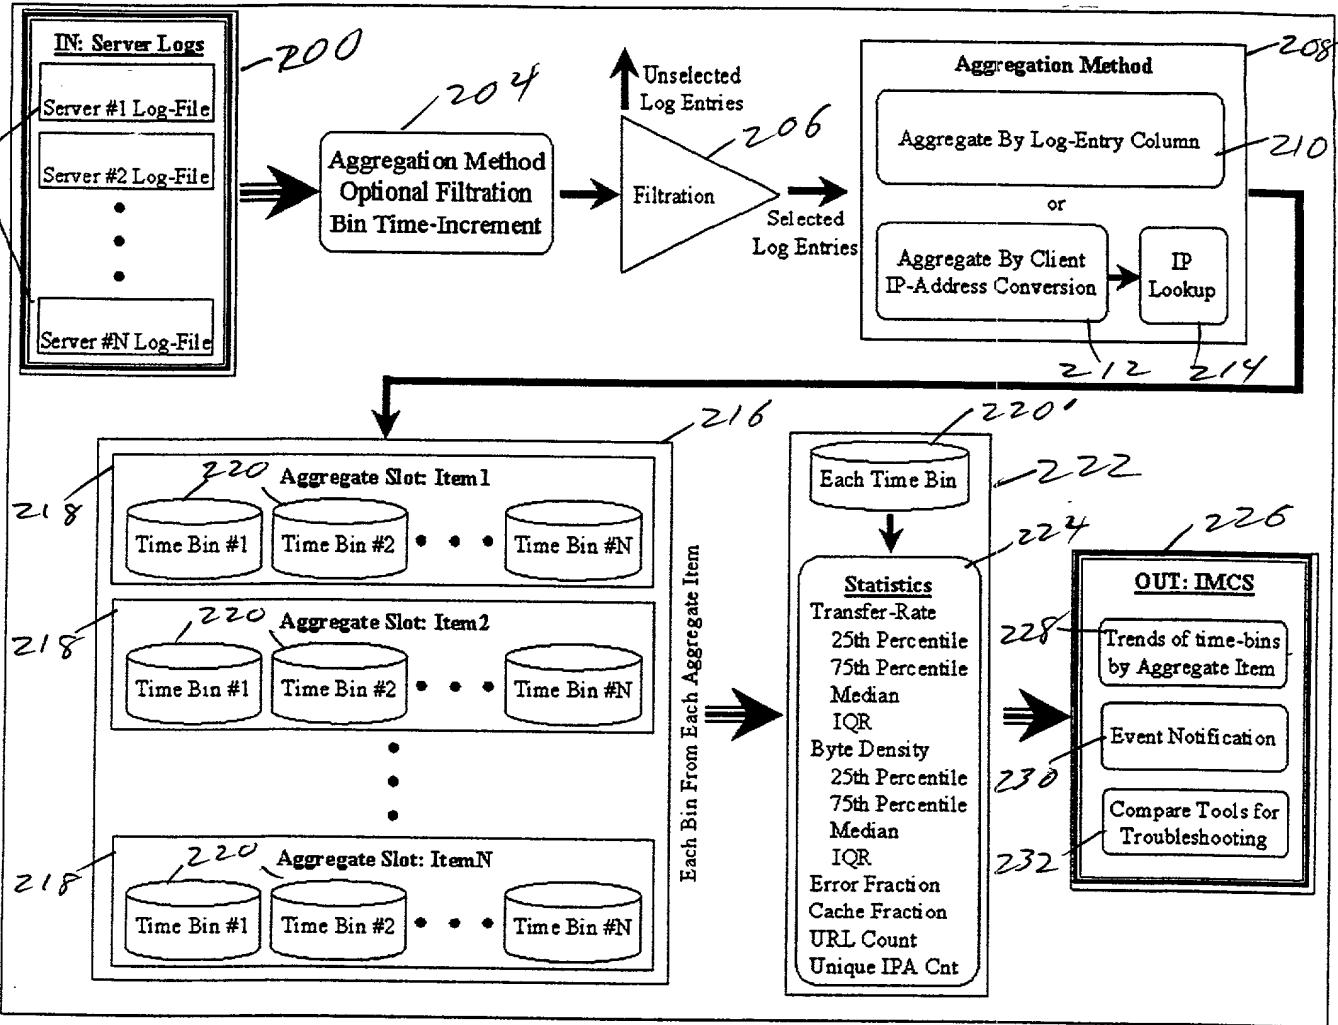
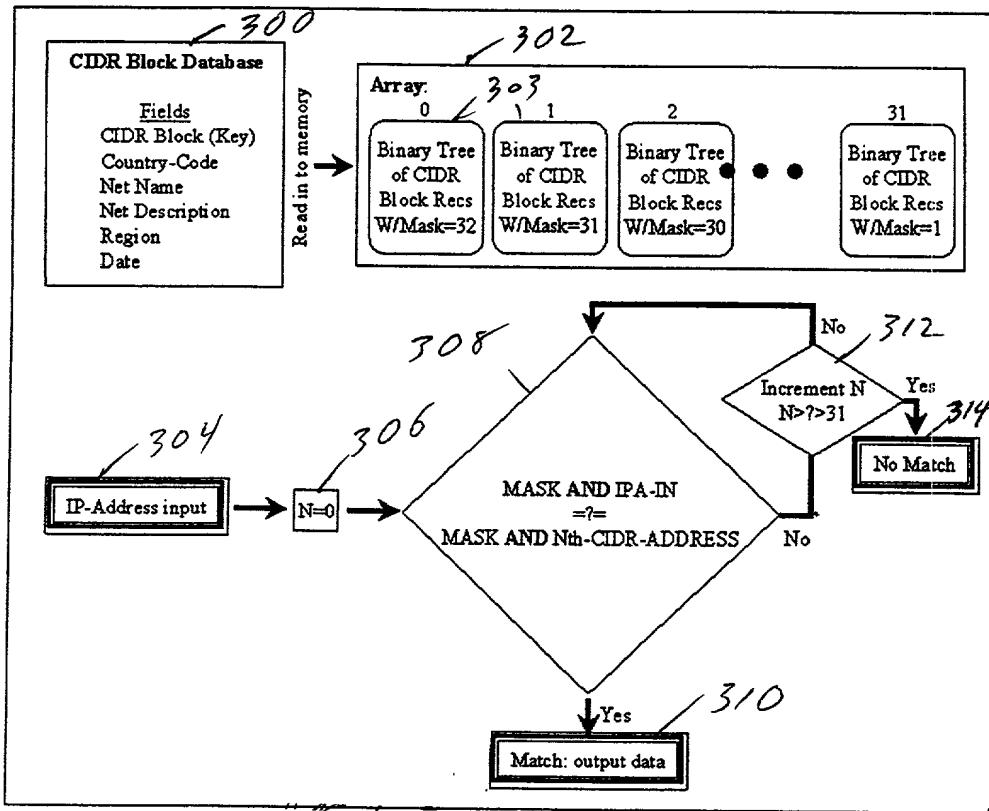


FIG. 2



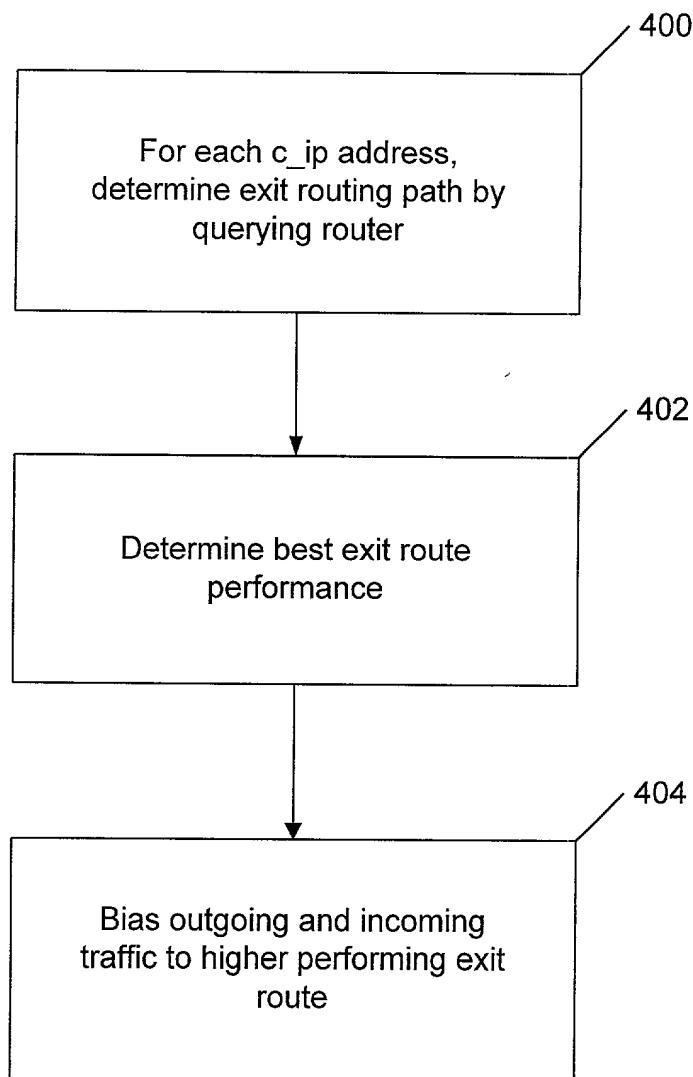


FIG. 4